

# Three-dimensional extinction mapping using Gaussian random fields

S. E. Sale and J. Magorrian

*Rudolf Peierls Centre for Theoretical Physics, Keble Road, Oxford OX1 3NP, UK*

Received ....., Accepted .....

## ABSTRACT

We present a scheme for using stellar catalogues to map the three-dimensional distributions of extinction and dust within our Galaxy. Extinction is modelled as a Gaussian random field, whose covariance function is set by a simple physical model of the ISM that assumes a Kolmogorov-like power spectrum of turbulent fluctuations. As extinction is modelled as a random field, the spatial resolution of the resulting maps is set naturally by the data available; there is no need to impose any spatial binning. We verify the validity of our scheme by testing it on simulated extinction fields and show that its precision is significantly improved over previous dust-mapping efforts. The approach we describe here can make use of any photometric, spectroscopic or astrometric data; it is not limited to any particular survey. Consequently, it can be applied to a wide range of data from both existing and future surveys.

**Key words:** ISM: dust, extinction – methods: statistical

## 1 INTRODUCTION

The extinction of starlight by interstellar dust is a nuisance for many astronomers. It is particularly troublesome in our own Galaxy: our understanding of the Galactic disc is hampered by in-plane dust, a frustrating situation given that the disc contains the bulk of the Galaxy’s stellar mass and is where the most interesting dynamical phenomena occur. There are already many stellar surveys that contain a wealth of still-undigested data about the structure and dynamics of the Galaxy, including those covering the whole sky (2MASS), targeted at the Galactic Plane (e.g. IPHAS Drew et al. 2005) or looking out of the plane (e.g. RAVE, SDSS). Many more, such as Gaia and Pan-STARRS and VPHAS+ (Drew et al. 2014), are underway. Any attempt to make sense of the catalogues produced by these surveys inevitably involves constructing model Galaxies and comparing the predictions from these models to the observed catalogues. For such a comparison to be meaningful, it is essential that the models’ predictions take account of the extinguishing effects of the Galaxy’s three-dimensional distribution of dust.

The well-known dust maps of, e.g., Schlegel, Finkbeiner & Davis (1998) or the Planck Collaboration et al. (2013) are not sufficient for this purpose, because they provide only the column density of dust integrated along each line of sight and have nothing to say about how the dust is distributed with distance. Nevertheless, these two-dimensional maps do provide important boundary conditions that any three-dimensional map must satisfy. There has been an increased focus on 3d mapping over the last several years. Marshall et al. (2006) compared 2MASS photometry to the Besançon Galactic model to produce a map of the extinction in much of the

Galactic plane. Majewski, Zasowski & Nidever (2011) presented the Rayleigh–Jeans colour excess method, which relies on the use of infrared data. Berry et al. (2012) estimated the distance and extinction to stars from SDSS and 2MASS, treating each star individually, before then considering the mean extinction of stars in spatial bins. Hanson & Bailer-Jones (2014), building upon Bailer-Jones (2011), proceeded along similar lines, though with a more sophisticated Bayesian method to estimate the distance and extinction to each star. A method which employs hierarchical Bayesian models to simultaneously estimate the reddening–distance relationship along a line of sight, and the properties of the stars that trace it, was presented in Sale (2012), superseding an earlier method by Sale et al. (2009). This has been applied by Sale et al. (2014) to construct a 3D extinction map of the Northern Galactic plane from the IPHAS DR2 catalogue (Barentsen et al. 2014). Green et al. (2014) adopted a similar approach, though the algorithmic details differ.

None of these methods properly account for the small-scale, fractal structure of the ISM; most assume that the dust column does not vary on scales of  $\sim 100$  pc. This assumption is fundamentally flawed, however: imaging studies show that the dust distribution has structure on sub-parsec scales (e.g. Figure 1 of Barentsen et al. 2011), with interstellar scintillation work demonstrating that the ISM has structure on scales as small as  $\sim 100$  km (e.g. Spangler & Gwinn 1990). Ignoring this substructure not only results in a loss of detail, but also leads to a potential bias in the resulting dust maps. Sale (2012) included a simple treatment for variations of extinction on small spatial scales, but Sale et al. (2014) found that the limitations of this treatment produced the dominant contribution to the uncertainty in the extinction to a given position, which can be significant at shorter distances.

An additional fact that is largely ignored by the preceding methods is that neighbouring patches of the ISM are not independent; turbulent processes fed by a variety of energy sources (Elmegreen & Scalo 2004) mix the gaseous components of the ISM. As dust is a passenger in the gas flows it too is mixed, causing its density to be correlated over the scales on which turbulence operates. Consequently, it is possible to constrain more tightly the density of components of the ISM by building these expected large-scale correlations into the mapping algorithm. This was recognised by Vergely et al. (2001) who applied geophysical mapping techniques from Tarantola & Valette (1982) to mapping the interstellar NaI and HII density. Subsequently, Vergely et al. (2010) and Lallement et al. (2014) have applied the same approach to mapping extinction in the local ISM. A similar approach was used by Pichon et al. (2001) to map the Lyman- $\alpha$  forest.

In this paper we present a method for constructing three-dimensional extinction maps based on modelling the dust density as a random field. Our method allows extinction to vary on all spatial scales by building in a simple physical model of the turbulence that produces structure in the ISM. The paper is organised as follows. Section 2 describes our model for the spatial distribution of dust. In Section 3 we tackle the general dust-mapping problem using an extension of the hierarchical Bayesian approach introduced by Sale (2012). We identify the posterior probability distributions that are most important for applications, and explain how they can be constructed from observations. To illustrate the application of the method, in Section 4 we use a simple two-dimensional dust-mapping problem. Section 5 compares our scheme to other methods and in section 6 we identify future work.

## 2 MODEL FOR THE SPATIAL DISTRIBUTION OF DUST

### 2.1 Motivation

No region of the ISM lives in isolation: the properties of the ISM at a given location are correlated with those of other positions, both in the immediate locality and at greater distances. These correlations are a product of the turbulent processes which shape the ISM over a range of spatial scales. We would like to establish a statistical framework to describe the distributions of dust and extinction, including a physical treatment of the correlations that exist as a product of turbulence.

#### 2.1.1 Spatial scaling

Kolmogorov (1941) proposed a now classic theory of turbulence, whereby energy is introduced to a medium at large scales. These processes cause the formation of large structures. Subsequently energy “cascades” down to smaller length scales, across an inertial range, producing structures of a corresponding size as it flows. Eventually, at small scales, viscosity becomes dominant and the energy is dissipated. In this picture dust is essentially a passenger, being dragged along by the gaseous components of the ISM which are themselves experiencing turbulence. However, as such it acts as a tracer of the gas distribution and so exhibits a similar density structure on the scales which are governed by turbulence.

One feature of Kolmogorov turbulence is that the density power spectrum of the medium takes the power-law form

$$|\bar{\rho}(\mathbf{k})|^2 \propto |\mathbf{k}|^{-\gamma}, \quad (1)$$

where  $\mathbf{k}$  represents the (three-dimensional) wavenumber and the

exponent  $\gamma = 11/3$  across the inertial range (e.g., Armstrong, Rickett & Spangler 1995). Observations confirm that this form is valid for the electron density over a wide range of scales in the diffuse ISM (Armstrong, Rickett & Spangler 1995; Chepurnov & Lazarian 2010). Similarly, HI surveys (e.g., Lazarian & Pogosyan 2000) and MHD simulations also obtain power-law power spectra, although these often display different slopes (e.g., Dickey et al. 2001; Padoan et al. 2004); the inclusion of magnetic fields (both in reality and MHD simulations) and the fact that the real ISM is not incompressible breaks a number of the assumptions made by Kolmogorov, giving rise to this discrepancy. Therefore, in the following we describe turbulence as being “Kolmogorov-like” if the power spectrum of its density fluctuations takes the power-law form (1), even if the power-law index  $\gamma \neq 11/3$ .

Lazarian & Pogosyan (2000), amongst others, demonstrate that the power spectra of projections of the density field, which include extinction, should follow a broken power-law. Specifically, the power spectrum slope changes at a scale corresponding to the transition between thin and thick screens of ISM. Maps of integrated Galactic dust emission show power-law angular power spectra (Schlegel, Finkbeiner & Davis 1998; Kiss et al. 2003; Roy et al. 2010) and similar spectra are also retrieved from 2D maps of extinction (Melbourne & Guhathakurta 2004; Brunt 2010).

#### 2.1.2 Distribution of density fluctuations at fixed position

Both simulations (e.g., Ostriker, Stone & Gammie 2001) and analytical arguments (e.g., Nordlund & Padoan 1999) indicate that, under a wide variety of conditions, the log density  $\log \rho(\mathbf{r})$  is approximately normally distributed, with a variance that depends on the local Mach number. But, note that Hopkins (2013) suggests an alternative distribution that may provide a better fit.

However, observations do not provide a direct probe of dust density, and so, for this paper, we are not immediately concerned with its probability distribution. Instead we view projections of the density field, such as extinction, which is itself directly related to the column density. For an observer located at the origin, the monochromatic<sup>1</sup> extinction to a point at distance  $s$  having galactic coordinates  $(l, b)$  is

$$A(l, b, s) \equiv \int_0^s \kappa \rho(l, b, s') ds', \quad (2)$$

where  $\kappa$  is the opacity and, for clarity, we have suppressed the wavelength dependence of both  $A$  and  $\kappa$  and any spatial dependence of  $\kappa$ . Clearly, the relationship between the statistical properties of  $A$  are easy to determine from those of  $\rho$ . The relationship between the properties of  $\log \rho$  and  $\log A$  is not so straightforward, however.

Nevertheless, Ostriker, Stone & Gammie (2001) reason that  $\log A$  should have an approximately Gaussian distribution. Their argument is as follows. Assume that the ISM along each line of sight is subject to many independent regions of compression or rarefaction. Each such region increases/decreases the density of the ISM by some fraction in the part of the sightline affected. Therefore it affects the projected density, along the entire line of sight, by another fraction, which is much closer to one. Consequently, the logarithms of these fractions are small. As the extinction contrast is the product of these fractions, it follows that the log extinction

<sup>1</sup> Following McCall (2004), Sale et al. (2009) and Bailer-Jones (2011), we consider only monochromatic extinctions in this paper, as they depend only on the dust column and not the properties of the observed star.

contrast is the sum of the logarithms. Given that the compression or rarefaction events are assumed to be independent, it follows that the log extinction contrast should follow a Gaussian distribution and so  $\log A$  should be normally distributed. This is borne out by both observations (Lombardi, Alves & Lada 2006; Kainulainen et al. 2009; Froebrich & Rowles 2010; Alves, Lombardi & Lada 2014) and simulations (Ostriker, Stone & Gammie 2001; Vázquez-Semadeni & García 2001), with only a small departure at the highest extinctions that occurs in regions of strong star formation.

## 2.2 A model for the statistical properties of the extinction distribution

The fundamental assumption of our method is that the logarithm of the extinction can be modelled as a semi-stationary Gaussian random field. To set the statistical properties of this field we use a model for the dust density that is based on the considerations above. Before describing our model in detail we first recall some basic properties of random fields in general, and Gaussian and semi-stationary random fields in particular.

### 2.2.1 Gaussian random fields

By a random field  $f(\mathbf{r})$  we mean a random process defined on three-dimensional space. We use  $f$  to refer both to the random process itself and to a particular function drawn from the random process. Given a set of  $N$  points,  $\mathbf{r}_1, \dots, \mathbf{r}_N$ , there is then a well-defined joint probability density function  $p(\mathbf{f})$  for the values  $\mathbf{f} \equiv (f(\mathbf{r}_1), \dots, f(\mathbf{r}_N))$  of  $f$  at those points. The first few moments of this pdf are a convenient (but incomplete) summary of the underlying random field. The most important are the mean and covariance, defined via

$$\begin{aligned}\bar{f}(\mathbf{r}) &\equiv \langle f(\mathbf{r}) \rangle, \\ \text{Cov}(f(\mathbf{r}_1), f(\mathbf{r}_2)) &\equiv \langle (f(\mathbf{r}_1) - \bar{f}(\mathbf{r}_1))(f(\mathbf{r}_2) - \bar{f}(\mathbf{r}_2)) \rangle \\ &= \langle f(\mathbf{r}_1)f(\mathbf{r}_2) \rangle - \bar{f}(\mathbf{r}_1)\bar{f}(\mathbf{r}_2),\end{aligned}\quad (3)$$

where  $\langle f(\mathbf{r}) \rangle$  denotes the value of  $f(\mathbf{r})$  averaged over many realizations of  $f$ . The quantity  $\langle f(\mathbf{r}_1)f(\mathbf{r}_2) \rangle$  that appears in the expression for the covariance is the correlation function.

A random field  $f$  is *stationary* if it is invariant under spatial translations. Clearly the expectation value of a stationary field is independent of position,  $\bar{f}(\mathbf{r}) = \text{constant}$ , and the covariance functions  $\text{Cov}[f(\mathbf{r}_i), f(\mathbf{r}_j)] = C(|\mathbf{r}_i - \mathbf{r}_j|)$  depends only on the scalar distance between points. By the Wiener-Khinchin theorem, the power spectrum  $|\tilde{f}(\mathbf{k})|^2$  and correlation function of a stationary field are a Fourier pair. That is,

$$\langle f(\mathbf{r})f(\mathbf{r} + \Delta\mathbf{r}) \rangle = \int d^3\mathbf{k} e^{i\mathbf{k} \cdot \Delta\mathbf{r}} |\tilde{f}(\mathbf{k})|^2, \quad (4)$$

where

$$\tilde{f}(\mathbf{k}) \equiv \frac{1}{(2\pi)^{3/2}} \int d^3\mathbf{r} e^{-i\mathbf{k} \cdot \mathbf{r}} f(\mathbf{r}) \quad (5)$$

is the Fourier transform of  $f(\mathbf{r})$ .

If we modulate a stationary field  $f(\mathbf{r})$  by multiplying it by a slowly varying function  $a(\mathbf{r})$ , the resulting  $g(\mathbf{r}) \equiv a(\mathbf{r})f(\mathbf{r})$  is an example of a *semi-stationary* process (Priestley 1965). Then, instead of the Wiener-Khinchin theorem, we have that

$$\langle g(\mathbf{r}_1)g(\mathbf{r}_2) \rangle \simeq a(\mathbf{r}_1)a(\mathbf{r}_2) \int d^3\mathbf{k} e^{i\mathbf{k} \cdot (\mathbf{r}_2 - \mathbf{r}_1)} |\tilde{f}(\mathbf{k})|^2. \quad (6)$$

This is a good approximation when the Fourier transform  $\tilde{a}(\mathbf{k})$  of the modulating function has an absolute maximum at  $\mathbf{k} = 0$ .

Finally,  $f(\mathbf{r})$  is a *Gaussian random field* if, for any choice of  $N$  points  $\mathbf{r}_1, \dots, \mathbf{r}_N$ , the joint pdf  $p(\mathbf{f})$  is a multivariate normal. That is,

$$p(\mathbf{f}) d\mathbf{f} = \frac{1}{(2\pi)^{3N/2} |\Sigma|^{N/2}} \exp \left[ -\frac{1}{2} (\mathbf{f} - \bar{\mathbf{f}})^T \Sigma^{-1} (\mathbf{f} - \bar{\mathbf{f}}) \right] d\mathbf{f}, \quad (7)$$

in which the mean values  $\bar{\mathbf{f}} \equiv (\bar{f}(\mathbf{r}_1), \dots, \bar{f}(\mathbf{r}_N))$  and the covariance matrix  $\Sigma_{ij} \equiv \text{Cov}[f(\mathbf{r}_i), f(\mathbf{r}_j)]$  depend only the location of the points  $\{\mathbf{r}_n\}$ . Such a field is completely specified by its mean  $\bar{f}(\mathbf{r})$  and covariance  $\text{Cov}[f(\mathbf{r}_i), f(\mathbf{r}_j)]$  functions. As the correlation function is related to the mean and covariance through

$$\langle f(\mathbf{r})f(\mathbf{r} + \Delta\mathbf{r}) \rangle = \bar{f}^2 + C(|\Delta\mathbf{r}|), \quad (8)$$

it follows that a stationary Gaussian random field is completely determined by its power spectrum  $|\tilde{f}(\mathbf{k})|^2$ . If  $f = \log A$  is a Gaussian random field, then  $A(\mathbf{r}) = e^{f(\mathbf{r})}$  is *not* a Gaussian random field, as the joint probability density of  $A(\mathbf{r}_1), \dots, A(\mathbf{r}_N)$  does not satisfy the condition (7). Instead, any field  $A(\mathbf{r})$  for which  $\log A$  is a Gaussian random field is called a *lognormal* random field.

### 2.2.2 Length scales

On the largest scales we do not expect the dust density to be a stationary process; clearly there are large regions where one expects the dust density to be relatively high (e.g., along spiral arms) and other regions where it will be low. We can, however, split our problem into two regimes by spatial scale: on larger scales the mean function of dust density will vary slowly with position, describing features such as spiral arms; whilst on smaller scales the turbulent ISM is resolved with a Gaussian random field. The question then becomes at which scale should the regimes be split and why?

We can address this by considering the power spectrum of dust density, which is generally thought to follow a power law across an inertial range between inner and outer limiting scales. The outer scale is determined by the dominant means of energy injection into the medium, with reduced power on larger scales. It is believed that the injection of energy into the ISM is dominated by supernovae, which produce features on spatial scales of the order of 100 pc (Elmegreen & Scalo 2004; MacLow 2004). Many other processes inject energy at larger or smaller scales, but are thought to add less energy than supernovae. Although the outer scale has not been identified in observations of Galactic dust, Haverkorn et al. (2008) and Chepurnov & Lazarian (2010) find that in interarm regions the outer scale of fluctuations in the warm ionised and neutral medium is on the order of 100 pc. Meanwhile, the power spectra of dust emission in M33 (Combes et al. 2012) is broken at a similar scale. Schlegel, Finkbeiner & Davis (1998) found no clear break in the power spectrum of Galactic dust emission, but we note that this is most likely a result of the projection of the 3D Galaxy onto a 2D image.

At small scales the onset of viscosity dominance reduces power and the ISM becomes essentially smooth. The inner scale is believed to be rather small. In the ionised interstellar medium it has been measured to be on the order of 100 km (Spangler & Gwinn 1990). There exist no observations of interstellar dust which probe such fine scales, but it is clear that it does exhibit structure on scales significantly smaller than a parsec (e.g. di Francesco et al. 2010; Robitaille, Joncas & Miville-Deschênes 2014). Therefore, the inner scale of turbulence can be safely ignored in this context, as we operate on significantly coarser scales.

### 2.2.3 Covariance function of $\rho$

Our model for the dust density is that  $\rho$  is a semi-stationary random field, with slowly varying mean  $m_\rho(\mathbf{r})$  and standard deviation  $S_\rho(\mathbf{r})$ . For example, this model can be obtained by expressing the logarithm of the density as

$$\log \rho(\mathbf{r}) = a_\rho(\mathbf{r}) + b_\rho(\mathbf{r})z(\mathbf{r}), \quad (9)$$

where  $z(\mathbf{r})$  is a stationary random field having zero mean and unit variance. The relationship between  $(a_\rho(\mathbf{r}), b_\rho(\mathbf{r}))$  and  $(m_\rho(\mathbf{r}), S_\rho(\mathbf{r}))$  depends on the distribution of  $z(\mathbf{r})$ . If we assume that  $z(\mathbf{r})$  is Gaussian, then

$$m_\rho(\mathbf{r}) = e^{a_\rho(\mathbf{r}) + b_\rho^2(\mathbf{r})/2}, \quad (10)$$

$$S_\rho^2(\mathbf{r}) = (e^{b_\rho^2(\mathbf{r})} - 1)e^{2a_\rho(\mathbf{r}) + b_\rho^2(\mathbf{r})}. \quad (11)$$

In Appendix B we show that this assumption produces a  $\log A$  distribution that is very close to Gaussian. We emphasise, however, that our method assumes only that  $\log A$  is a Gaussian random field; the details of the  $\rho$  field that produces this lognormal extinction distribution are not important, provided only that  $\rho$  is some semi-stationary random field, described by some mean  $m_\rho(\mathbf{r})$ , variance  $S_\rho(\mathbf{r})$  and power spectrum.

Following the considerations outlined in Section 2.1.1, we want the density power spectrum to have the Kolomogorov-like form (1). But, as noted in Section 2.2.2, this scale-free form cannot hold for all wavenumbers  $|\mathbf{k}|$ . So, we take the power spectrum of  $\rho$  to be proportional to the function

$$\Upsilon_{\gamma,\Omega}(\mathbf{k}, \mathcal{L}_o) \equiv R(\gamma, \Omega, \mathcal{L}_o) \frac{(|\mathbf{k}| \mathcal{L}_o)^{2\Omega}}{[1 + (|\mathbf{k}| \mathcal{L}_o)^2]^{\frac{7}{2} + \Omega}}, \quad (12)$$

in which the normalization constant  $R(\gamma, \Omega, \mathcal{L}_o)$  is given by

$$\frac{1}{R(\gamma, \Omega, \mathcal{L}_o)} = 4\pi \int_0^\infty d|\mathbf{k}| \frac{|\mathbf{k}|^2 (|\mathbf{k}| \mathcal{L}_o)^{2\Omega}}{[1 + (|\mathbf{k}| \mathcal{L}_o)^2]^{\frac{7}{2} + \Omega}}, \quad (13)$$

so that taking  $|\tilde{\rho}(\mathbf{k})|^2 = \Upsilon_{\gamma,\Omega}(\mathbf{k}, \mathcal{L}_o)$  results in a density field with unit variance. We then scale the spectrum by  $S_\rho^2(\mathbf{r})$  to account for variations in the standard deviation of the field and add a  $\delta$ -function at the origin, scaled by  $m_\rho^2(\mathbf{r})$ , to encode variations in the mean density. This spectrum has a Kolmogorov-like  $|\mathbf{k}|^{-\gamma}$  shape within the characteristic scalelength  $\mathcal{L}_o$ . On larger scales the spectrum tapers smoothly to zero, the shape of the taper being set by the parameter  $\Omega$ . By tapering the spectrum, structures on scales greater than  $\mathcal{L}_o$  are described only by variations in  $m_\rho(\mathbf{r})$  and  $S_\rho(\mathbf{r})$  and are unaffected by the random field  $z(\mathbf{r})$ . This gives rise to the split in scales discussed in section 2.2.2. We do not adjust (12) to account for the inner scale of turbulence, as this is significantly smaller than any of the scales we can probe in the ISM.

We note that, as indicated by equation (4), setting the power spectrum of  $\rho(\mathbf{r})$  is equivalent to setting its covariance function, since the two are a Fourier pair. However, setting the power spectrum is more convenient, because it offers a straightforward way of employing a physically motivated Kolmogorov-like spectrum. Fig. 1 compares power spectra of different forms and the covariance function pairs of the power spectra considered.

### 2.2.4 Covariance function of $\log A$

We model  $\log A$  as a Gaussian random field whose mean and covariance functions depend on the density function above. The mean is

straightforward to calculate from the function  $m_\rho(\mathbf{r})$ , given the covariance function of  $\log A$ . The relationship between the covariance of  $\log A$  and the density is less simple and so we relegate the details to Appendix A. Equation (A8) gives the covariance function of  $\log A$  which depends on the power spectrum (12) we choose for the dust density as well as on  $m_\rho(\mathbf{r})$  and  $S_\rho(\mathbf{r})$ , but does not depend on the form of the distribution of  $\log \rho$ .

## 3 CONSTRUCTING MAPS

The problem we address in this paper is that of inferring the three-dimensional distribution of extinction from some data. Typically these data will come in the form of a catalogue or catalogues of observations of some set of stars, whether photometric, spectroscopic, astrometric, or some combination of the three. For example, one could employ 2MASS photometry on its own, or use 2MASS combined with IPHAS photometry, or 2MASS plus a spectroscopic survey, such as SEGUE, RAVE or APOGEE. Subsequently, one might want to refine the maps by including parallaxes from Gaia as they become available.


We split the problem into three tasks. First, given some catalogue of observations of stars, we would like to estimate the extinction towards each of the stars in the catalogue: as discussed in Sale (2012), a properly constructed, coherent extinction map also helps us determine other properties of these stars – such as their distance or masses – more precisely than would be possible if the stars were studied separately. Second, from this catalogue we would like to infer the extinction to arbitrary positions in space. For example, we could be interested in the extinction to stars or other objects that were not included in the catalogue used to construct the map. Finally, we seek to characterise the distribution of extinction and, by extension, dust. We are interested in how they are distributed on large scales as well as the detailed statistical behaviour of dust which informs us about the turbulent processes that shape the ISM.

In the following we introduce a barrage of labels to denote various parameters and modelling assumptions. These are summarised in Table 1.

### 3.1 Probability of a map

We use  $\Theta$  to stand for the set of parameters of the functions  $m_\rho(\mathbf{r})$  and  $S_\rho(\mathbf{r})$  that describe the large-scale distribution of extinction. These parameters could be as simple as a scale height and scale length describing an exponential disc. A more sophisticated model might use three-dimensional splines, in which case  $\Theta$  would contain the positions of the spline points, the values of  $m_\rho(\mathbf{r})$  and  $S_\rho(\mathbf{r})$  at each point, and any prior assumptions about the covariances among the values. We use  $\zeta$  to denote our model for the small-scale properties of the ISM (such as  $\gamma$ ,  $\mathcal{L}_o$  and  $\Omega$ ) and  $\alpha$  for our model for *how* stars, subject to some extinction, would appear in observations. For example, a simple model for the latter may include isochrones, synthetic stellar atmospheres and a model for the wavelength dependence of extinction.

We are given a catalogue of  $N$  stars. Let  $(\tilde{l}_n, \tilde{b}_n)$  be the observed Galactic coordinates and let  $\tilde{y}_n$  be the full set of all other observations (observed parallax, apparent magnitudes, etc) of the  $n^{\text{th}}$  star in the catalogue. We use the tilde to indicate that these are observations and not the true values. Each star has some unknown distance,  $s_n$ , and extinction,  $A_n$ , along with a set of other intrinsic parameters (mass, age, metallicity and so on) that we denote by  $x_n$ . We use vectors  $(\tilde{\mathbf{l}}, \tilde{\mathbf{b}}) = \{\{\tilde{l}_n\}, \{\tilde{b}_n\}\}$  and  $\tilde{\mathbf{y}} = \{\tilde{y}_n\}$  to denote the



**Figure 1.** A comparison of the power spectra (left) and normalised covariance functions (right) with four different forms of truncation at large scales. The black solid line shows the untruncated form, the red dashed line shows sudden truncation at  $k_0$ , whilst the  $Y_{11/3,0}$  and  $Y_{11/3,1}$  forms are shown with blue dot-dashed and green dashed lines respectively.

$(\tilde{l}_n, \tilde{b}_n)$	the $(l, b)$ coordinates observed for the $n^{\text{th}}$ star in the catalogue.
$\tilde{y}_n$	all other directly observed quantities (e.g., broad-band fluxes, trigonometric parallax) for the $n^{\text{th}}$ star.
$s_n$	the distance to the $n^{\text{th}}$ star in the catalogue.
$A_n$	the extinction to the $n^{\text{th}}$ star in the catalogue.
$x_n$	all other intrinsic properties of the $n^{\text{th}}$ star (mass, metallicity, age, etc).
$\Theta$	parameters that set the large-scale distribution of extinction, $m_p(\mathbf{r})$ , $S_p(\mathbf{r})$
$\eta$	any assumptions made about $\Theta$ (e.g., an assumed functional form for $m_p(\mathbf{r})$ , $S_p(\mathbf{r})$ )
$\zeta$	the model for the small scale distribution of $\log \rho$ , including $(\gamma, L_0, \Omega)$
$\beta$	the background Galaxy model, including prior on position, metallicity, age, etc.
$\alpha$	denotes model for mapping $(s_n, A_n)$ to predicted observables, $\tilde{y}_n$ .

**Table 1.** A list of the parameters and assumptions used in our model.

full set of observations in the catalogue. Similarly,  $\mathbf{s} = \{s_n\}$  and  $\mathbf{A} = \{A_n\}$  refer to the full set of distances and extinctions to each star in the catalogue, while  $(\mathbf{l}, \mathbf{b}) = (\{l_n\}, \{b_n\})$  denote the true Galactic coordinates of the stars in the catalogue.

We assume a prior model,  $\beta$ , for the joint probability distribution of the parameters  $(l_n, b_n, s_n, x_n)$  describing each star's three-dimensional position, age, metallicity, and so on. We treat all but the position of the star and its extinction as nuisance parameters and use the method of Sale (2012), or equivalently Hanson & Bailer-Jones (2014) or Green et al. (2014), to marginalise  $x_n$  to obtain the likelihood  $p(\tilde{l}_n, \tilde{b}_n, \tilde{y}_n | l_n, b_n, s_n, A_n, \beta, \alpha)$  of  $(l_n, b_n, s_n, A_n)$  for each star given the assumed physical model  $\alpha$  for the dust extinction and the background galaxy model  $\beta$ . Examples of these marginalised likelihoods can be seen in Figure 6 of Green et al. (2014).

Using this likelihood, the posterior distribution

$$\begin{aligned}
 & p(\Theta, \mathbf{l}, \mathbf{b}, \mathbf{s}, \mathbf{A} | \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \tilde{\mathbf{y}}, \zeta, \alpha, \beta, \eta) \\
 & \propto p(\mathbf{A} | \mathbf{l}, \mathbf{b}, \mathbf{s}, \Theta, \zeta) p(\Theta | \eta) \\
 & \times \prod_n p(\tilde{l}_n, \tilde{b}_n, \tilde{y}_n | l_n, b_n, s_n, A_n, \alpha, \beta) p(l_n, b_n, s_n, | \beta).
 \end{aligned} \tag{14}$$

Here  $p(\mathbf{A} | \mathbf{l}, \mathbf{b}, \mathbf{s}, \Theta, \zeta)$  is a multivariate Gaussian on  $\log \mathbf{A}$ , the covariance matrix of which is set using the covariance function  $\text{Cov}[\log A(l_1, b_1, s_1), \log A(l_2, b_2, s_2)]$  given by equation (A8), which in turn depends on an assumed dust power spectrum as in equation (12). The mean vector for  $\log \mathbf{A}$  is found by first integrating the mean density along the line of sight to each star to find the expected extinction to each star. By combining these and the covariance matrix it is then possible to find the expectation of log-extinction for all the stars.

We assume that our observations of the on-sky position of each

star are independent of the other observations of that star:

$$\begin{aligned}
 & p(\tilde{l}_n, \tilde{b}_n, \tilde{y}_n | l_n, b_n, s_n, A_n, \alpha, \beta) \\
 & = p(\tilde{l}_n, \tilde{b}_n | l_n, b_n, \alpha) p(\tilde{y}_n | l_n, b_n, s_n, A_n, \alpha, \beta).
 \end{aligned} \tag{15}$$

Typically the uncertainties in  $(\tilde{l}_n, \tilde{b}_n)$  are negligible and we may take

$$p(\tilde{l}_n, \tilde{b}_n | l_n, b_n, \alpha) = \delta(\tilde{l}_n - l_n) \delta(\tilde{b}_n - b_n). \tag{16}$$

Marginalising  $(\mathbf{l}, \mathbf{b})$ , equation (14) simplifies to

$$\begin{aligned}
 & p(\Theta, \mathbf{s}, \mathbf{A} | \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \tilde{\mathbf{y}}, \zeta, \alpha, \beta, \eta) \\
 & = p(\mathbf{A} | \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \mathbf{s}, \Theta, \zeta) p(\Theta | \eta) \\
 & \times \prod_n p(\tilde{y}_n | \tilde{l}_n, \tilde{b}_n, s_n, A_n, \alpha, \beta) p(s_n | \tilde{l}_n, \tilde{b}_n, \beta).
 \end{aligned} \tag{17}$$

In  $p(\Theta | \eta)$  we include any assumptions we wish to make about the parameters that define the large-scale structure of extinction. For example, a simple model for the dust density might include a characteristic scale height among the parameters  $\Theta$ . Then the prior assumed for this scale height would be incorporated into  $p(\Theta | \eta)$ . Finally,  $p(s_n | \tilde{l}_n, \tilde{b}_n, \beta)$  is a prior on the distance to the stars based on the background model  $\beta$ .

### 3.2 Extinction and distance to a single star in the catalogue

If we are interested in a particular star, we can obtain the marginal posterior for only this star by marginalising over the parameters of

the other stars and  $\Theta$ :

$$\begin{aligned}
 & p(s_1, A_1 | \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \tilde{\mathbf{y}}, \zeta, \alpha, \beta, \eta) \\
 &= \int p(\Theta, \mathbf{s}, \mathbf{A} | \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \tilde{\mathbf{y}}, \zeta, \alpha, \beta, \eta) d\mathbf{s}_{2 \dots N} d\mathbf{A}_{2 \dots N} d\Theta \\
 &= p(\tilde{\mathbf{y}}_1 | \tilde{\mathbf{l}}_1, \tilde{\mathbf{b}}_1, s_1, A_1, \alpha, \beta) p(s_1 | \tilde{\mathbf{l}}_1, \tilde{\mathbf{b}}_1, \beta) \\
 &\quad \times \int d\Theta \int d\mathbf{s}_{2 \dots N} \int d\mathbf{A}_{2 \dots N} \left[ p(\mathbf{A} | \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \mathbf{s}, \Theta, \zeta) p(\Theta | \eta) \right. \\
 &\quad \left. \times \prod_{n=2}^N p(\tilde{\mathbf{y}}_n | \tilde{\mathbf{l}}_n, \tilde{\mathbf{b}}_n, s_n, A_n, \alpha, \beta) p(s_n | \tilde{\mathbf{l}}_n, \tilde{\mathbf{b}}_n, \beta) \right].
 \end{aligned} \tag{18}$$

Notice how all the other stars influence the posterior of the star we are interested in. If we took this star to be independent of all the others the posterior would reduce to the product of the first two terms on the right hand side of equation (19). By using the extra information given by the observations of the other stars, however, it is possible to obtain a more precise estimate of the distance and extinction to the star than if it had been considered in isolation.<sup>2</sup> It might appear that such an integration above would be expensive to perform, but it is naturally obtained from the MCMC algorithms we use to sample the posterior (17); a sample from the marginalised posterior on  $s_1$  and  $A_1$  is found by only considering these variables in the MCMC chain.

### 3.3 Inferring extinction to arbitrary position(s)

Another question one might ask is, what is the extinction to an arbitrary position or positions? For example, one might want to know the extinction to a non-stellar object or to a star that did not appear in the original catalogue. We can obtain the extinction  $A^*$  to any position  $(l^*, b^*, s^*)$  from the posterior predictive distribution,

$$\begin{aligned}
 & p(A^* | l^*, b^*, s^*, \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \tilde{\mathbf{y}}, \zeta, \alpha, \beta, \eta) \\
 &= \int d\mathbf{A} d\mathbf{s} d\Theta \\
 &\quad p(A^* | l^*, b^*, s^*, \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \mathbf{s}, \mathbf{A}, \Theta, \zeta) p(\Theta, \mathbf{s}, \mathbf{A} | \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \tilde{\mathbf{y}}, \zeta, \alpha, \beta, \eta).
 \end{aligned} \tag{20}$$

This marginalisation is most straightforwardly carried out by using an MCMC algorithm to explore the posterior distribution  $p(\Theta, \mathbf{s}, \mathbf{A} | \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \tilde{\mathbf{y}}, \zeta, \alpha, \beta, \eta)$  and summing up the samples to obtain a realisation of  $p(A^* | l^*, b^*, s^*, \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \mathbf{s}, \mathbf{A}, \Theta, \zeta)$ . The posterior is obtained directly from our assumption that  $\log A$  is Gaussian random field. This lognormal assumption means that finding the joint probability distribution of the extinction to multiple points in space is an obvious, straightforward generalisation of (20)

### 3.4 General structure of dust

One can also obtain the marginal posterior of the large-scale structure of dust,  $p(\Theta | \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \tilde{\mathbf{y}}, \zeta, \alpha, \beta)$ , by marginalising  $\mathbf{s}$  and  $\mathbf{A}$  from equation (17). As in section 3.2, this can be easily achieved if an MCMC algorithm has been employed.

We postpone to a future paper the question of how to infer the parameters  $(\gamma, \mathcal{L}, \Omega)$  that govern the model  $\zeta$  for the small-scale physics. For now we simply assume that  $\gamma = 11/3$ ,  $\mathcal{L}_0 = 100$  pc and  $\Omega = 1$ .

<sup>2</sup> Of course, if our model for the covariance of  $\log A$  is wrong, our inferred  $A_n$  and  $s_n$  – however precise – will be biased.

## 4 EXAMPLES

We now consider some examples to illustrate how to use the formalism above to construct extinction maps within a narrow, two-dimensional wedge of a model galaxy. We restrict ourselves to two dimensions only because it is then straightforward to present the results on paper; there is nothing fundamentally different in our scheme between this simple two-dimensional wedge and the full three-dimensional case. As we work in two dimensions we can describe a star's position with its distance  $s_n$  and Galactic longitude  $l_n$  only, enabling us to drop the Galactic latitude  $b_n$ .

To focus attention on the dust model, the model galaxy  $\beta$  we consider here is deliberately crude: there is a single population of stars, the spatial density of which is uniform within the wedge, so that

$$p(s_n | l_n, \beta) \propto s_n. \tag{21}$$

The dust density model  $\Theta$  is slightly more realistic. The function that sets the mean density of the dust distribution falls off exponentially with distance from the observer:

$$\bar{\rho}(l, s) \equiv m_p(l, s) = \rho_0 \exp\left(-\frac{s}{s_l}\right). \tag{22}$$

This is intended to mimic the situation in which one is modelling observations towards the Galactic anticentre  $(l, b) = (180^\circ, 0^\circ)$ . We take the ratio between the standard deviation and mean of dust density to be constant  $S_p(l, s)/m_p(l, s) = \sqrt{e-1}$ , where  $e$  is the base of natural logarithms. Therefore our model  $\Theta$  describing the large-scale dust density has only two free parameters,  $\rho_0$  and  $s_l$ . In the following we adopt a uniform prior on  $\log \rho_0$  and  $\log s_l$ .

We avoid specifying an explicit model  $\alpha$  for the extinguishing effects of dust on the observables  $\tilde{\mathbf{y}}_n$  by taking  $\tilde{\mathbf{y}}_n = (\log \tilde{s}_n, \log \tilde{A}_n)$  and assigning observational uncertainties to the “observed” (log) distance,  $\log \tilde{s}_n$ , and “observed” (log) extinction,  $\log \tilde{A}_n$ , to each star.

To construct our simulated galaxy catalogue we first use the procedure described in appendix B1 to generate a three-dimensional realisation of the dust density with power spectrum proportional to (12), which are modulated by the mean density  $m_p(s)$  given by equation (22) having scale length  $s_l = 2$  kpc and asymptotic extinction

$$A_\infty \equiv \kappa \rho_0 s_l = 5.0. \tag{23}$$

Then we sprinkle 200 stars on a wedge of length 10 kpc and opening angle  $30'$  within this volume. We read off the true distances and extinctions  $y_{n,\text{true}} = (s_{n,\text{true}}, A_{n,\text{true}})$  to each star, then scatter each by an amount consistent with the assumed observational uncertainty to obtain the “observed” values  $\tilde{\mathbf{y}}_n$ .

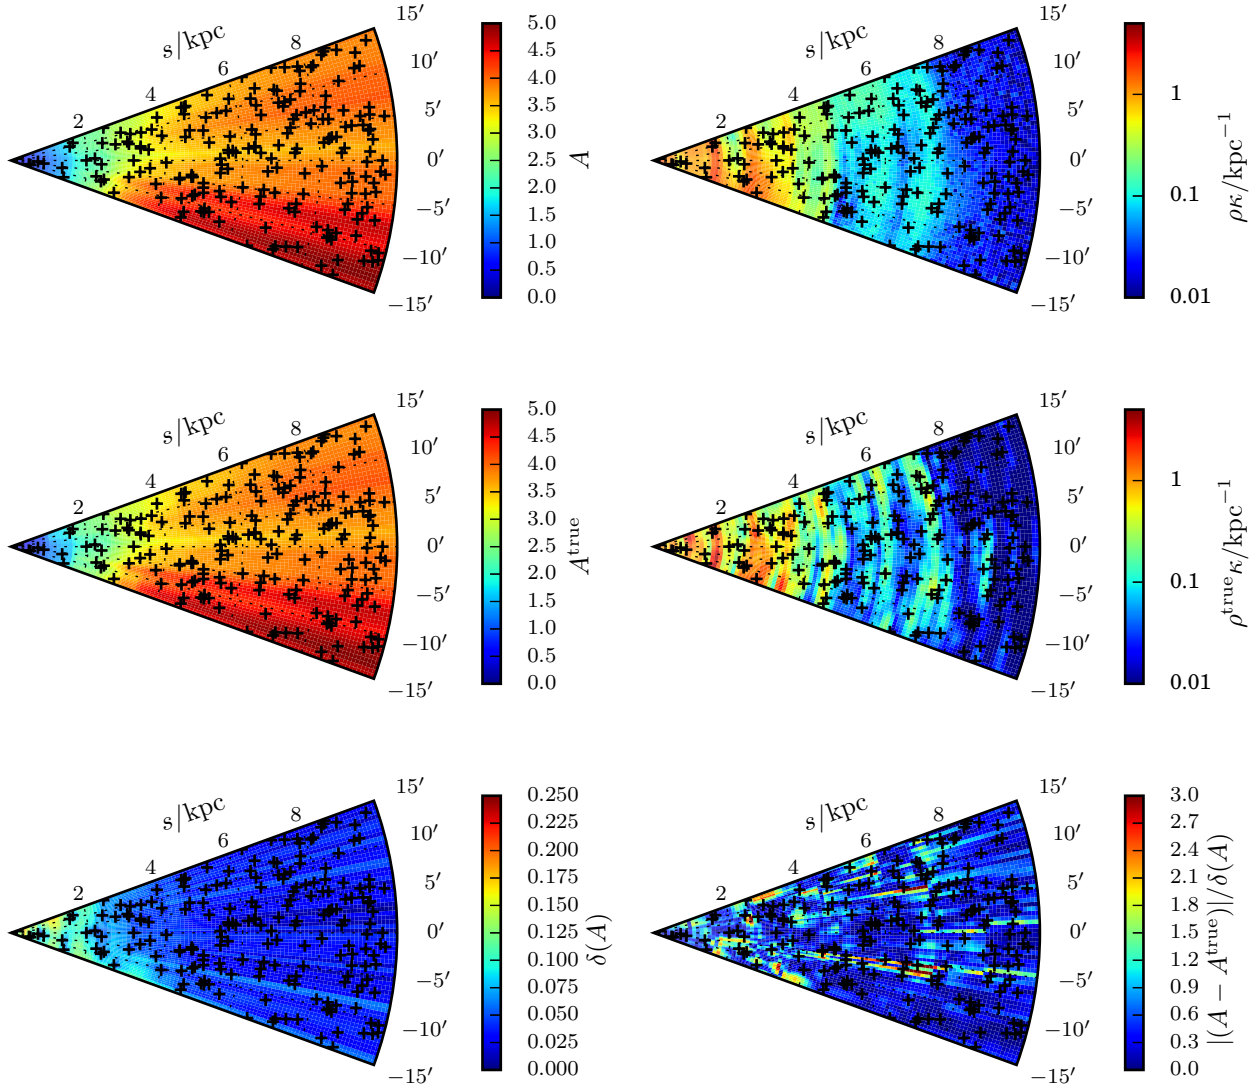
### 4.1 Perfect Data

We first consider the case where we possess perfect data, meaning that we know the parameters of the stars with absolute precision. Then

$$\begin{aligned}
 & p(\tilde{\mathbf{y}} | \tilde{\mathbf{l}}, \mathbf{s}, \mathbf{A}, \alpha) p(\mathbf{s} | \tilde{\mathbf{l}}, \beta) \\
 &= \prod_n \delta(\log s_n - \log s_{n,\text{true}}) \delta(\log A_n - \log A_{n,\text{true}}).
 \end{aligned} \tag{24}$$

The only uncertain parameters are those describing the mean extinction,  $\Theta = (a_0, s_l)$ , and, instead of considering the full posterior described by equation (17), we need only consider

$$p(\Theta | \tilde{\mathbf{y}}, \tilde{\mathbf{l}}, \zeta, \eta) \propto p(\mathbf{A}_{\text{true}} | \tilde{\mathbf{l}}, \mathbf{s}_{\text{true}}, \Theta, \zeta) p(\Theta | \eta), \tag{25}$$



**Figure 4.** A series of maps demonstrating the results our method achieves from ‘perfect’ simulated data. We plot the posterior expectation of extinction (top left) and density (top right), the true extinction (middle left) and density (middle right), the width of a 68% credible interval on extinction as a measure of uncertainty (bottom left) and the residual between the posterior expectation of extinction and true extinction normalised by the uncertainty. On all the plots crosses indicate the positions of the stars ‘observed’.

where  $p(\mathbf{A}_{\text{true}} | \tilde{\mathbf{l}}, \mathbf{s}_{\text{true}}, \Theta, \zeta)$  is a Gaussian in  $\log(\mathbf{A}_{\text{true}})$ , with a covariance matrix set using the covariance function described in appendix A. This posterior can be estimated using a simple Markov chain Monte Carlo (MCMC) algorithm to sample the members of  $\Theta$ . We run an MCMC algorithm with a Metropolis-Hastings updater for  $\Theta$  for 20000 iterations, discarding the first 5000 iterations in the chain as burn-in. Both the overall length of the chain and the length of burn in are rather conservative: the autocorrelation length of the chain is less than 40 iterations and the chain typically converges within the first thousand iterations.

The marginal posteriors on  $s_l$  and  $A_\infty$  for an example realization are shown in Figures 2 and 3. Subsequently we continue by calculating the posterior predictive distribution  $p(A^* | l^*, s^*, \tilde{\mathbf{l}}, \tilde{\mathbf{y}}, \alpha, \beta, \zeta, \eta)$ . This requires marginalising over  $\Theta$ , which is achieved simply by iterating over the different members

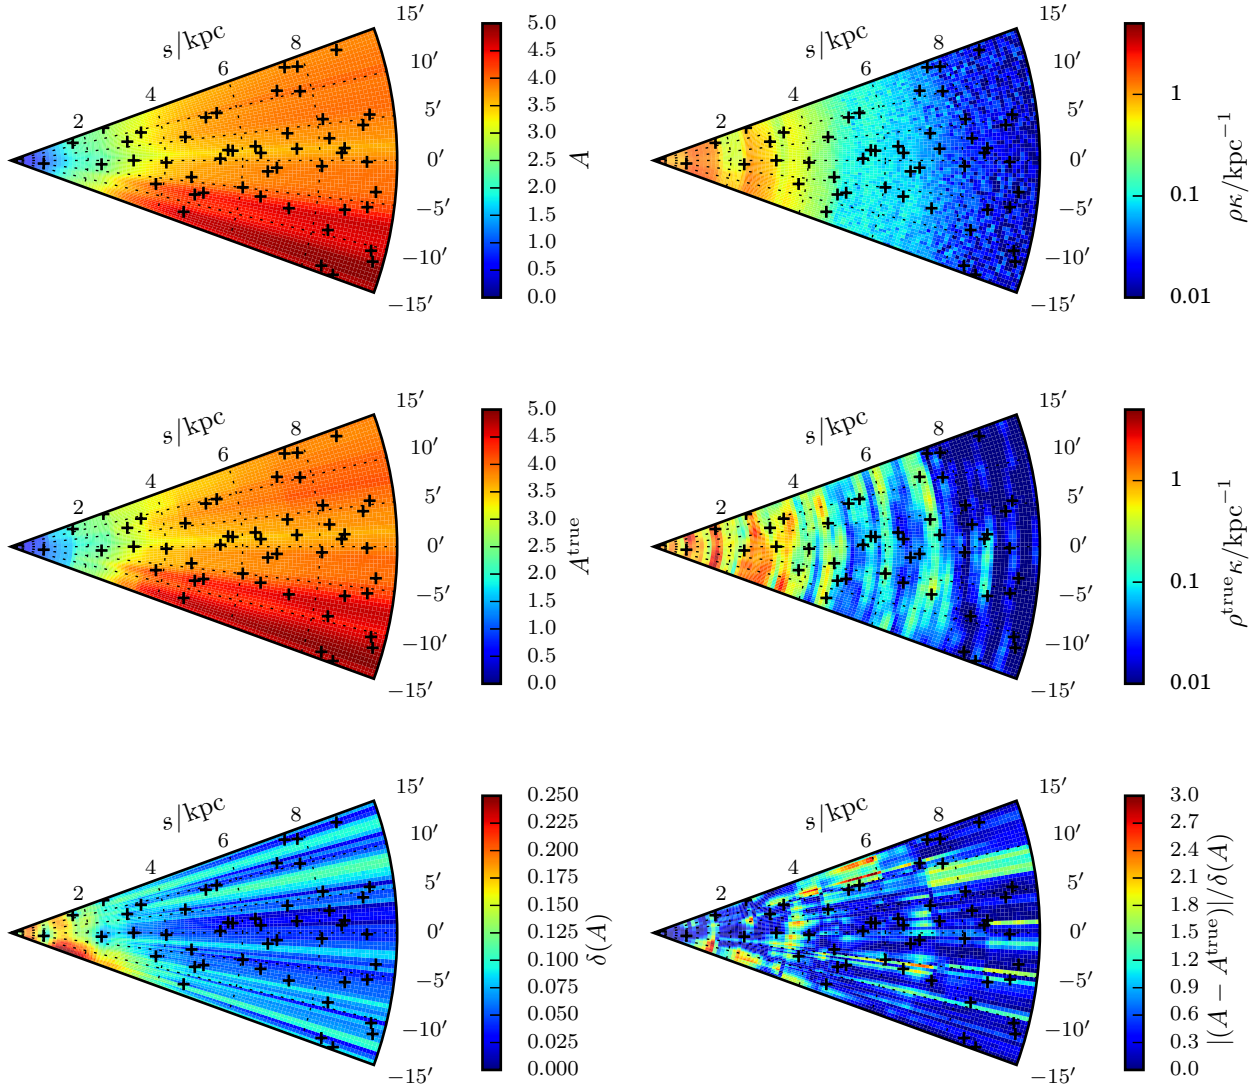
of the MCMC chain for  $\Theta$ . The extinction map produced in this way is shown in Figure 4.

The resolution of the map is fundamentally limited by the distribution of stars. To demonstrate this we remove 150 stars from the catalogue used in this section and reanalyse. Figure 5 shows the extinction map obtained. There is a clear loss of detail, particularly in the density map and a small increase in uncertainty. None the less the larger extinction features are still found.


## 4.2 Imperfect data

Next we consider the effects of observational uncertainties by adding random errors to the true  $(\mathbf{s}, \mathbf{A})$  as follows. For each star






**Figure 5.** As with Fig 4, but only employing 50 of the original 200 stars.



**Figure 2.** The marginal posterior distribution of dust scale length  $s_l$  found from perfect two-dimensional data. The true value of  $s_l$  used to simulate the data is  $s_l = 2$  kpc.



**Figure 3.** The marginal posterior distribution of asymptotic extinction  $A_\infty$  found from perfect two-dimensional data. The true value of  $A_\infty$  used to simulate the data is  $A_\infty = 5$ .



we draw  $\Delta y_n$  from a Gaussian having zero mean and covariance

$$\Sigma_n = \begin{pmatrix} \frac{\Sigma_{\log s, n}}{\tau_n \sqrt{\Sigma_{\log s, n} \Sigma_{\log A, n}}} & \tau_n \sqrt{\Sigma_{\log s, n} \Sigma_{\log A, n}} \\ \tau_n \sqrt{\Sigma_{\log s, n} \Sigma_{\log A, n}} & \Sigma_{\log A, n} \end{pmatrix}. \quad (26)$$

The variable  $\tau_n$  is a random number drawn from the uniform distribution  $[-1, 1]$  and is included to model correlations between the observations of  $s_n$  and  $A_n$  for a single star. Note that we do not include any star–star correlations:  $\Delta y_n$  is independent of  $\Delta y_m$  for  $m \neq n$ . We model  $\Sigma_{\log A, n}$  in a such a way that it grows as the extinction and/or distance to the star increases. This is done by first giving each star a ‘true’ apparent magnitude in the  $V$  band,

$$V_n = 2 + 5 \log_{10}(s_{n, \text{true}}/10\text{pc}) + A_{n, \text{true}}. \quad (27)$$

Uncertainties on extinction are then assumed to grow exponentially with apparent magnitude, with variance

$$\Sigma_{A, n} = e^{1.4 + 1.5(V_n - 23)} + 0.0025, \quad (28)$$

from which it is straightforward to find  $\Sigma_{\log A, n}$ . We further set  $\Sigma_{\log s, n} = 4\Sigma_{\log A, n}$ . We set the uncertainties in this manner to imitate the growth of photometric uncertainties. The specific scaling is applied to approximately match median uncertainties found from IPHAS data in Sale et al. (2014). With this scaling, the rms scatter in  $(\log s, \log A)$  varies from  $\sim (0.06, 0.03)$  for a star at 2 kpc to  $\sim (0.44, 0.22)$  for one at 10 kpc. We add this  $\Delta y_n$  to  $y_{n, \text{true}}$  to obtain our “observations”,

$$\tilde{y}_n = \begin{pmatrix} \log \tilde{s}_n \\ \log \tilde{A}_n \end{pmatrix}, \quad (29)$$

where  $\tilde{s}_n$  is the “observed” distance and  $\tilde{A}_n$  the “observed” extinction to a star.

To model these observational uncertainties in our map-making machinery we replace the the Dirac deltas in (24) by the Gaussians

$$p(\tilde{\mathbf{y}}|\tilde{\mathbf{l}}, \mathbf{s}, \mathbf{A}, \alpha, \beta)p(\mathbf{s}|\tilde{\mathbf{l}}, \beta) \\ \propto \prod_n \frac{1}{2\pi|\Sigma_n|^{1/2}} \exp \left[ -\frac{1}{2} \left( \tilde{y}_n - \begin{pmatrix} \log s_n \\ \log A_n \end{pmatrix} \right)^T \Sigma_n^{-1} \left( \tilde{y}_n - \begin{pmatrix} \log s_n \\ \log A_n \end{pmatrix} \right) \right] \quad (30)$$


With this it is straightforward to marginalise out  $\mathbf{A}$  from the posterior equation (17): for given  $\mathbf{s}$ , the integral of  $p(\tilde{\mathbf{y}}|\tilde{\mathbf{l}}, \mathbf{s}, \mathbf{A}, \alpha, \beta)p(\mathbf{A}|\tilde{\mathbf{l}}, \mathbf{s}, \Theta, \zeta)$  with respect to  $\log \mathbf{A}$  is simply a convolution of Gaussians. Then we have that

$$p(\Theta, \mathbf{s}|\tilde{\mathbf{l}}, \tilde{\mathbf{y}}, \alpha, \beta, \zeta, \eta) \\ \propto p(\tilde{\mathbf{A}}|\tilde{\mathbf{s}}, \tilde{\mathbf{l}}, \mathbf{s}, \Theta, \alpha, \zeta)p(\tilde{\mathbf{s}}|\mathbf{s}, \alpha)p(\mathbf{s}|\tilde{\mathbf{l}}, \beta)p(\Theta|\eta), \quad (31)$$

in which the first two factors are Gaussians in  $\log \tilde{\mathbf{A}}$  and  $\log \tilde{\mathbf{s}}$  respectively. The covariance matrix for  $p(\tilde{\mathbf{A}}|\tilde{\mathbf{s}}, \tilde{\mathbf{l}}, \mathbf{s}, \Theta, \alpha, \zeta)$  is equal to that of  $p(\mathbf{A}|\tilde{\mathbf{l}}, \mathbf{s}, \Theta, \zeta)$  plus the  $\Sigma_{\log A, n}$  along the diagonal, whilst the mean vector is unchanged.

An obvious MCMC scheme for this posterior is to alternate updates of  $\Theta$  with updates to  $\mathbf{s}$ : this is an example of a Metropolis-within-Gibbs MCMC scheme (Tierney 1994). We continue to employ a simple Metropolis-Hastings updater for  $\Theta$ , whilst all members of  $\mathbf{s}$  are updated simultaneously with another Metropolis-Hastings updater, as updating each member in sequence would require recalculating  $p(\tilde{\mathbf{A}}|\tilde{\mathbf{s}}, \tilde{\mathbf{l}}, \mathbf{s}, \Theta, \alpha, \zeta)$  for each member, which is computationally prohibitive.

The justification for marginalising over  $\mathbf{A}$  follows from practical experience with this MCMC algorithm: the autocorrelation time



**Figure 7.** A histogram of the posterior expectation of extinction to 5 kpc at an angle of  $0^\circ$  minus the ‘true’ extinction to this position and divided by the standard deviation of the posterior. The measured mean of the data is 0.004 and the standard deviation 1.008. Over plotted with a red line is a unit Gaussian.

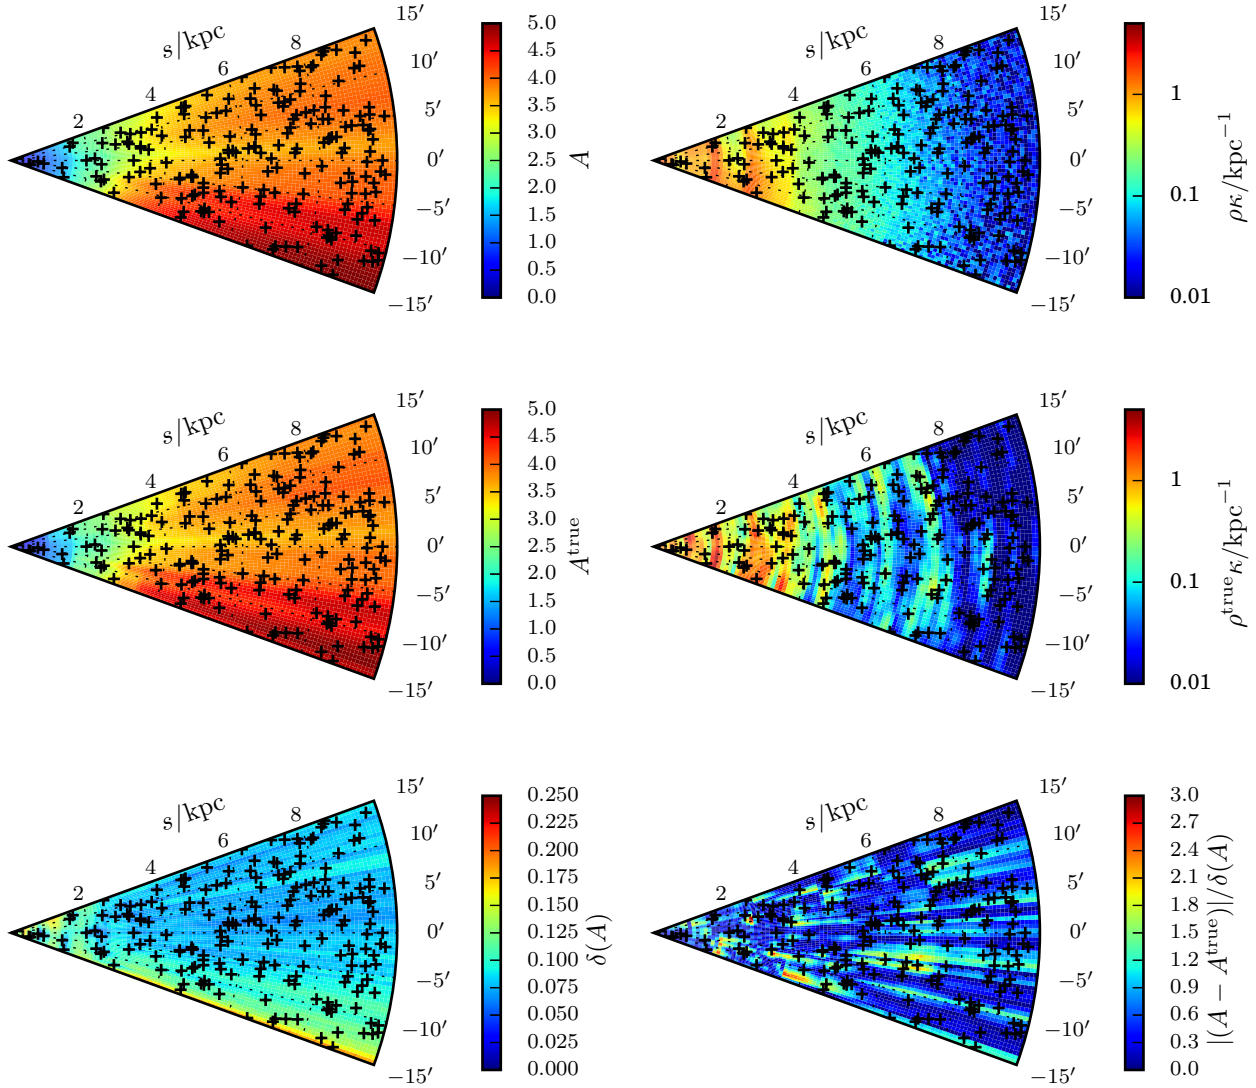
in the marginalised case is  $\sim 10$  times less than if  $\mathbf{A}$  were also sampled. If we desire, once we have used MCMC to estimate the posterior in equation (31), we can then extend this to find the posterior on  $\mathbf{A}$  using a similar process to that described in section 3.3.

Figure 6 shows an example of the extinction and dust density maps reconstructed from our simulated catalogues. The uncertainty in the extinction and density maps has increased relative to the case of perfect data, with some corresponding loss of detail, particularly in the density map. However, the extinction map still largely reproduces the input.

To get a more rigorous handle on the reliability of our method we have simulated 5000 independent extinction fields and star catalogues, adding random observational errors to each. We then analysed each as before. To show an example of how we have used these to appraise the quality of our results we consider extinction along the line of sight  $l = 0^\circ$  to a distance of 5 kpc. For each realization we have recorded the difference between the posterior expectation value of extinction to this position and the ‘true’ value. For comparison, we have also recorded the standard deviation of the posterior distribution, which serves to quantify the uncertainty in the posterior estimate. Fig. 7 shows the distribution of the ratio of these two quantities – that is, the difference of the estimate from the true value divided by the uncertainty in the estimate. The mean of the distribution shown is 0.004, confirming that the posterior expectation of extinction is a good estimator of the true extinction. Additionally, the measured standard deviation of the plotted distribution is 1.008, suggesting that the standard deviation of the posterior is a reliable estimate of the uncertainty.

## 5 DISCUSSION

Sale et al. (2014) presented a 3D extinction map of the northern Galactic Plane, based on a method from Sale (2012). That method proceeds by dividing the sky into many regions approximately  $10'$  in size and then estimating both mean extinction and ‘differential extinction’ – the unresolved variation in extinction – as a function of distance. One of the most striking features of the 3D extinction map it produces is that uncertainty in the extinction to a given point is dominated by the effects of unresolved substructure. Consequently it was apparent that a method that is more directly geared towards dealing with the fractal nature of the ISM was necessary.



**Figure 6.** As with Fig 4, but using data with simulated errors and uncertainties.


The method we present here fits that requirement: a physical model for the small-scale structure of the ISM is centre stage.

To demonstrate the improvements of our method with respect to Sale (2012), we use the problems of estimating extinction both to particular stars and to arbitrary points in space. We show marginal posteriors for both these problems in Figs. 8 and 9, comparing the two methods. The difference between the two methods is best appreciated by comparing the widths of the posterior distributions: those obtained from our new method are considerably narrower and therefore more precise than those found by the older approach. The improved precision stems directly from the more sophisticated treatment of small-scale structure in the ISM. We also note that although the posterior expectations of extinction in both Figs. 8 and 9 do not align exactly with the true values, there is no substantial bias as evidenced by Fig. 7.


A significant undesirable feature present in the maps of Marshall et al. (2006) and Sale et al. (2014) are ‘fingers of god’ – radially extended, azimuthal discontinuities in maps of dust den-

sity. These artefacts appear because the methods of Sale (2012) and Marshall et al. (2006) treat separate sight lines independently. As our model for small scale structure of the ISM includes a spatial correlation kernel for dust density such features no longer occur.

Our decision to model the likelihood as a bivariate normal (30) on  $\{\log A_n, \log s_n\}$  was motivated primarily by the opportunity to perform the marginalisation of  $\mathbf{A}$  analytically. The lognormal dependence of this likelihood on distance is plausibly realistic, as it corresponds to a Gaussian likelihood on distance modulus. However, Green et al. (2014) show that when we take  $y_n$  to be real observables, such as a set of apparent magnitudes in different bands, then the likelihood  $p(\bar{y}_n | \bar{l}_n, \bar{b}_n, s_n, A_n, \alpha, \beta)$  typically takes a complicated form, reflecting the shape of stellar isochrones. Consequently a single Gaussian in  $\{\log s_n, \log A_n\}$  space is unlikely to always provide a good fit. This can be remedied by fitting a mixture of Gaussians to each likelihood; the generalisation of our method from a single Gaussian to a mixture of Gaussians is trivial in principle and requires only some modest extra bookkeeping in practice.



**Figure 8.** Estimated posterior probability distributions of extinction for a single star located at 5169 pc. The ‘true’ extinction to the star is  $A_0 = 3.72$ , as indicated by the vertical dashed blue line. In black is a histogram of samples from  $p(A_i|\tilde{l}, \tilde{y}, \zeta, \beta, \alpha, \eta)$ , i.e. the posterior obtained from the method we describe once the distance of this star, the distance and extinction of all other stars and the large scale structure of dust have been marginalised over. The red histogram represents a similar histogram obtained using the method of Sale (2012). Finally in green we have the posterior distribution obtained by considering the star in isolation, i.e. without using the observations of other stars. Formally this is  $p(A_i|\tilde{l}_i, \tilde{y}_i, \beta, \alpha)$ .



**Figure 9.** Estimated posterior probability distributions of extinction to a position at 5000 pc and an angle of  $0^\circ$ . The ‘true’ extinction to this point is  $A_0 = 3.52$ , as indicated by the vertical dashed blue line. The black histogram shows the probability distribution of  $p(A^*|l^*, s^*, \tilde{l}, \tilde{y}, \zeta, \alpha, \beta, \eta)$  estimated using the technique described in section 3.3. The red histogram shows a comparable estimate obtained using the method of Sale (2012).

The method we have described largely follows and builds upon that of Vergely et al. (2001). The two methods are fundamentally similar: both employ a Gaussian random field to describe  $\log A$ . However, in the models of Vergely et al. (2001), the large-scale ISM density is not allowed to vary. This worked well in their study, as they only studied a relatively local volume. Our method allows us to infer the large-scale distribution of dust, a necessary step as we seek to map extinction to much greater distances, where the distribution of dust is not well known.

Furthermore, we diverge when it comes to inferring the posterior. We use an MCMC based approach to sample from the posterior, whilst Vergely et al. (2001) use an iterative inversion algorithm from Tarantola & Valette (1982) to find the expectation of the posterior. The principal benefit of the Tarantola & Valette (1982) algorithm is speed: Vergely et al. (2001) note that they typ-

ically converge on their final solution within 10 iterations. However, this algorithm is only valid for solving an (almost) linear least squares problem. Consequently, the posterior, likelihood and priors can only follow Gaussian distributions. In contrast our MCMC based algorithm is able to deal with non-Gaussian pdfs. The method in Vergely et al. (2001) is also unable to easily deal with uncertainties on the distance to each star, requiring that they be absorbed into extinction estimates. This will be particularly troublesome if uncertainties on a star’s distance and extinction are correlated – as is generally the case – since it will result in a loss of precision.

Selection functions can have a particularly pathological impact on extinction mapping: a magnitude limited sample will preferentially include less extinguished stars, with the consequence that, if the selection function is not accounted for, the produced extinction map will be biased to lower extinctions (Sale 2012). As selection functions contribute a strongly non-Gaussian component to the posterior, Vergely et al. (2001) are unable to deal with them effectively. However, they avoid this problem by only employing relatively local catalogues which are at least approximately volume limited, a small proportion of available data. We have not directly approached selection functions in this paper: in all the simulations we performed it was assumed that all simulated stars were ‘observed’. However, we note that the method we present has the flexibility to properly deal with the selection functions which shape the catalogues we employ. We will consider this in more detail in a future paper.

Vergely et al. (2010) adopts a sum of two exponentials with different scales as their covariance function, whilst Lallement et al. (2014) employ a sum of a Gaussian and a sech term. They suggest that the two scales are indicative of the ISM being shaped by two similarly strong processes which occur on different scales. We note that the summed kernels they employ can be used to provide a reasonable approximation of the covariance function partners of members of our  $\Gamma_{\gamma, \Omega}$  family (12). Thus we suggest that the fact they find that they need to sum two kernels arises because the simple kernels they employ provide a poor approximation of a Kolmogorov-like covariance function, but a sum of two or more simple kernels provides a reasonable approximation. Therefore, there is no need to invoke a second process injecting energy into the ISM.

## 6 SUMMARY AND FUTURE WORK

The coming years will see vast quantities of astrophysical survey data become available from ongoing (e.g. Gaia, PAN-STARRS) and upcoming (e.g. LSST) surveys, which will join the large resources we already have from surveys such as 2MASS and SDSS. Given the ubiquity of extinction and its corresponding imprint on the data these surveys will furnish, a precise treatment of it is vitally important in order to fully exploit these data. Moreover, the distribution of extinction is shaped by a range of processes operating in the ISM, from the formation of spiral arms to the enrichment of the ISM by AGB stars. Consequently the study of extinction grants us a view on these processes.

Our statistical model includes a realistic physical description of the small-scale structure of the ISM which has generally been neglected when mapping extinction. We have assumed a fixed form for the covariance function for dust density in this paper, with the shape set by the outer scale length  $\mathcal{L}_o$  and the slope of the corresponding power spectrum  $\gamma$ . However, these values are not well known and may indeed vary across the Galaxy (Haverkorn et al.

2008). We intend to address this in a future paper, where we will extend the existing method to attempt to estimate these values from real observations.

We have yet to mention the wavelength dependence of reddening in this paper, or the possibility that it might vary with position. One possible approach would be to simply assume that the form of the reddening law to a star is independent of that to all other stars. Then it would be straightforward to marginalise over the form of the reddening law to obtain a likelihood for each star conditioned on the stars' distances and extinction, but not the form of the extinction law. However, such an approach is clearly suboptimal in the sense that we expect that the form of the reddening law to two stars to be correlated if these two stars are located near each other. Again we defer a deeper investigation of this issue to a subsequent paper.

We note that calculating  $p(\mathbf{A}|\tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \mathbf{s}, \Theta, \zeta)$  or  $p(\tilde{\mathbf{A}}|\tilde{\mathbf{s}}, \tilde{\mathbf{l}}, \tilde{\mathbf{b}}, \mathbf{s}, \Theta, \alpha, \zeta)$ , which follow multivariate lognormal distributions, is the most computationally expensive part of the MCMC algorithms employed in this paper as it involves solving for the inverse and determinant of the covariance matrix. Moreover, naive algorithms for finding the inverse and the determinant of a matrix require  $O(N^3)$  time, where  $N$  is the number of stars in the catalogue. In the future we intend to apply the method we have described to large catalogues. Clearly it will not be trivial to do so: we would have to deal with an extremely large covariance matrix, which, under a naive method would be impossible to invert on any reasonable time scale. However, as the stars in such a catalogue will be widely spread, the overwhelming majority of pairs of stars will be at least approximately independent of each other and so the covariance matrix will be very sparse. Therefore, it should be possible to leverage this considerable sparsity to make the method feasible on large scales.

We have demonstrated, using simulated data, that our method can successfully retrieve the distribution of extinction in a field and that it is several times more precise than the method of Sale (2012). Once the computational challenges have been overcome we will be in possession of a method for inferring three-dimensional extinction maps from large stellar catalogues, through the use of Gaussian random fields and a physical model for the small-scale fractal structure of dust. We will then be able to apply this method to the vast torrent of survey data to produce a beautifully precise and compelling three dimensional map of extinction.

## ACKNOWLEDGEMENTS

We would like to thank Janet Drew, Christophe Pichon, the Oxford Dynamics group and the anonymous referee for their contributions and suggestions that have improved this paper.

We acknowledge support from the United Kingdom Science Technology and Facilities Council (STFC, ST/K00106X/1).

## REFERENCES

Alves J., Lombardi M., Lada C. J., 2014, *A&A*, 565, A18  
 Armstrong J. W., Rickett B. J., Spangler S. R., 1995, *ApJ*, 443, 209  
 Bailer-Jones C. A. L., 2011, *MNRAS*, 411, 435  
 Barakat R., 1976, *J. Opt. Soc. America*, 66, 211  
 Barentsen G. et al., 2014, preprint (arXiv:1406.4862)  
 —, 2011, *MNRAS*, 415, 103

Berry M. et al., 2012, *ApJ*, 757, 166  
 Brunt C. M., 2010, *A&A*, 513, A67  
 Chepurinov A., Lazarian A., 2010, *ApJ*, 710, 853  
 Coles P., Jones B., 1991, *MNRAS*, 248, 1  
 Combes F. et al., 2012, *A&A*, 539, A67  
 di Francesco J. et al., 2010, *A&A*, 518, L91  
 Dickey J. M., McClure-Griffiths N. M., Stanimirović S., Gaensler B. M., Green A. J., 2001, *ApJ*, 561, 264  
 Drew J. E. et al., 2014, *MNRAS*, 440, 2036  
 —, 2005, *MNRAS*, 362, 753  
 Elmegreen B. G., 2002, *ApJ*, 564, 773  
 Elmegreen B. G., Scalo J., 2004, *ARA&A*, 42, 211  
 Fischera J., Dopita M. A., 2004, *ApJ*, 611, 919  
 Froebrich D., Rowles J., 2010, *MNRAS*, 406, 1350  
 Green G. M. et al., 2014, *ApJ*, 783, 114  
 Hanson R. J., Bailer-Jones C. A. L., 2014, *MNRAS*, 438, 2938  
 Haverkorn M., Brown J. C., Gaensler B. M., McClure-Griffiths N. M., 2008, *ApJ*, 680, 362  
 Hopkins P. F., 2013, *MNRAS*, 430, 1880  
 Ju N., 2002, *Journal of Computational Finance*, 5, 79  
 Kainulainen J., Beuther H., Henning T., Plume R., 2009, *A&A*, 508, L35  
 Kiss C., Ábrahám P., Klaas U., Lemke D., Héraudeau P., del Burgo C., Herbstmeier U., 2003, *A&A*, 399, 177  
 Kolmogorov A., 1941, *Akademiia Nauk SSSR Doklady*, 30, 301  
 Lallement R., Vergely J.-L., Valette B., Puspitarini L., Eyer L., Casagrande L., 2014, *A&A*, 561, A91  
 Lazarian A., Pogosyan D., 2000, *ApJ*, 537, 720  
 Lombardi M., Alves J., Lada C. J., 2006, *A&A*, 454, 781  
 MacLow M.-M., 2004, *AP&SS*, 289, 323  
 Majewski S. R., Zasowski G., Nidever D. L., 2011, *ApJ*, 739, 25  
 Marshall D. J., Robin A. C., Reylé C., Schultheis M., Picaud S., 2006, *A&A*, 453, 635  
 McCall M. L., 2004, *AJ*, 128, 2144  
 Mehta N. B., Wu J., Molisch A. F., Zhang J., 2007, *Wireless Communications, IEEE Transactions on*, 6, 2690  
 Melbourne J., Guhathakurta P., 2004, *AJ*, 128, 271  
 Nordlund Å. K., Padoan P., 1999, in *Interstellar Turbulence*, Franco J., Carraminana A., eds., p. 218  
 Ostriker E. C., Stone J. M., Gammie C. F., 2001, *ApJ*, 546, 980  
 Padoan P., Jimenez R., Juvela M., Nordlund Å., 2004, *ApJ*, 604, L49  
 Pichon C., Vergely J. L., Rollinde E., Colombi S., Petitjean P., 2001, *MNRAS*, 326, 597  
 Planck Collaboration et al., 2013, preprint (arXiv:1312.1300)  
 Priestley M. B., 1965, *Journal of the Royal Statistical Society. Series B (Methodological)*, 204  
 Robitaille J.-F., Joncas G., Miville-Deschênes M.-A., 2014, *MNRAS*, 440, 2726  
 Roy A. et al., 2010, *ApJ*, 708, 1611  
 Sale S. E., 2012, *MNRAS*, 427, 2119  
 Sale S. E. et al., 2014, *MNRAS*, 443, 2907  
 —, 2009, *MNRAS*, 392, 497  
 Saupe D., 1988, *The science of fractal images*, 71  
 Schlegel D. J., Finkbeiner D. P., Davis M., 1998, *ApJ*, 500, 525  
 Spangler S. R., Gwinn C. R., 1990, *ApJ*, 353, L29  
 Tarantola A., Valette B., 1982, *Reviews of Geophysics and Space Physics*, 20, 219  
 Tierney L., 1994, *Annals of Statistics*, 22, 1701  
 Vázquez-Semadeni E., Garcia N., 2001, *ApJ*, 557, 727  
 Vergely J.-L., Freire Ferrero R., Siebert A., Valette B., 2001, *A&A*, 366, 1016

Vergely J.-L., Valette B., Lallement R., Raimond S., 2010, A&A, 518, A31

## APPENDIX A: APPROXIMATING THE EXTINCTION COVARIANCE FUNCTION FROM THE DENSITY POWER SPECTRUM

In this appendix we obtain the covariance function of extinction  $A$  given the power spectrum for the dust density  $\rho$ . The latter is modelled as a semi-stationary random field (Priestley 1965) whose power spectrum is given a function  $\mathfrak{v}(\mathbf{k})$ , where this power spectrum is tapered at small  $|\mathbf{k}|$ . For example, we may choose to use a member of the family  $\Upsilon_{\gamma,\Omega}(\mathbf{k})$ , as given by equation (12). The power spectrum is then modulated by the effects of the slowly varying functions  $m_\rho(\mathbf{r})$  and  $S_\rho(\mathbf{r})$ . We use this covariance function of  $A$  to obtain an approximate expression for the covariance function of  $\log A$ .

We place an observer at the origin  $O$  of our coordinate system and consider the extinctions

$$A(\mathbf{s}_i) = \int_0^{s_i} \kappa \rho(s'_i \hat{\mathbf{s}}_i) ds'_i \quad (\text{A1})$$

to an arbitrary pair of points  $(\mathbf{s}_1, \mathbf{s}_2)$ , in which the  $\hat{\mathbf{s}}_i$  are unit vectors along the line of sight to each point and the scalars  $s_i$  measure the distances of each from the observer. The correlation function

$$\langle A(\mathbf{s}_1) A(\mathbf{s}_2) \rangle = \int_0^{s_1} ds'_1 \int_0^{s_2} ds'_2 \langle \kappa \rho(\mathbf{s}'_1) \kappa \rho(\mathbf{s}'_2) \rangle, \quad (\text{A2})$$

where we have introduced  $\mathbf{s}'_i \equiv s'_i \hat{\mathbf{s}}_i$  to keep notation reasonably uncluttered. As we assume that  $\rho$  (strictly,  $\kappa \rho$ ) is semi-stationary, we may use equation (6) to approximate this as

$$\langle A(\mathbf{s}_1) A(\mathbf{s}_2) \rangle \simeq \int_0^{s_1} ds'_1 \int_0^{s_2} ds'_2 \langle \kappa \rho(\mathbf{s}'_1) \rangle \langle \kappa \rho(\mathbf{s}'_2) \rangle \int d^3 \mathbf{k} e^{i\mathbf{k} \cdot (\mathbf{s}'_1 - \mathbf{s}'_2)} \mathfrak{v}(\mathbf{k}). \quad (\text{A3})$$

To simplify this further, we assume that the angle  $\Delta\theta$  between  $\hat{\mathbf{s}}_1$  and  $\hat{\mathbf{s}}_2$  is small, and choose a coordinate system in which  $\hat{\mathbf{s}}_1 = (\frac{1}{2}\Delta\theta, 0, 1)$  and  $\hat{\mathbf{s}}_2 = (-\frac{1}{2}\Delta\theta, 0, 1)$ . That is, the  $z$  axis of this coordinate system bisects the lines of sight to  $\mathbf{s}_1$  and  $\mathbf{s}_2$ . Making the approximation  $\mathfrak{v}(k_x, k_y, k_z) \simeq \mathfrak{v}(k_x, k_y, 0)$  and using the Fourier-space representation of the Dirac delta,

$$\delta(s) = \frac{1}{2\pi} \int_{-\infty}^{\infty} dk_z e^{ik_z s}, \quad (\text{A4})$$

we obtain

$$\langle A(\mathbf{s}_1) A(\mathbf{s}_2) \rangle \simeq 2\pi \int_0^{\min(s_1, s_2)} ds \langle \kappa \rho(s \hat{\mathbf{s}}_1) \rangle \langle \kappa \rho(s \hat{\mathbf{s}}_2) \rangle \int_{-\infty}^{\infty} dk_x \int_{-\infty}^{\infty} dk_y e^{\frac{i}{2} k_x s \Delta\theta} \mathfrak{v}(k_x, k_y, 0). \quad (\text{A5})$$

This can be rewritten as

$$\langle A(\mathbf{s}_1) A(\mathbf{s}_2) \rangle \simeq 2\pi \int_0^{\min(s_1, s_2)} ds \langle \kappa \rho(s \hat{\mathbf{s}}_1) \rangle \langle \kappa \rho(s \hat{\mathbf{s}}_2) \rangle P(s\Delta\theta), \quad (\text{A6})$$

where  $\Delta\theta = \cos^{-1}(\hat{\mathbf{s}}_1 \cdot \hat{\mathbf{s}}_2)$  and the function

$$P(x) \equiv \int_{-\infty}^{\infty} dk_x \int_{-\infty}^{\infty} dk_y e^{\frac{i}{2} k_x x} \mathfrak{v}(k_x, k_y, 0) \quad (\text{A7})$$

can easily be tabulated for a given choice of power spectrum  $\mathfrak{v}(\mathbf{k})$ .

We have made a number of assumptions to obtain this result. First, although we have not assumed any particular form for the power spectrum of turbulence, we do assume that it is isotropic and semi-stationary. Second, we have made the approximation  $\mathfrak{v}(k_x, k_y, k_z) \simeq \mathfrak{v}(k_x, k_y, 0)$  by taking a Taylor expansion of  $\mathfrak{v}(k_x, k_y, k_z)$  around  $k_z = 0$  and discarding all but the leading term, a valid approximation if our sightlines  $s_1, s_2$  are much longer than  $\mathcal{L}_0$ . If we wished we could include further terms in the Taylor expansion, particularly if we were dealing with short sightlines. Third, we have assumed that  $\Delta\theta \ll 1$ , but note that satisfying the condition that  $s_1, s_2 \gg \mathcal{L}_0$  ensures that one need not worry about the small angle assumption, as  $P(s\Delta\theta) \rightarrow 0$  when  $s\Delta\theta \gg \mathcal{L}_0$ .

Up until now, we have made no assumptions about the Gaussianity of the random fields  $\rho$  and  $A$ . If we now assume that  $\log A$  can be modelled as a Gaussian random field, we can follow Coles & Jones (1991) and use the result above to approximate the covariance function of  $\log A$  as

$$\text{Cov}[\log A(\mathbf{s}_1), \log A(\mathbf{s}_2)] \simeq \log \left[ \frac{\langle A(\mathbf{s}_1) A(\mathbf{s}_2) \rangle}{\langle A(\mathbf{s}_1) \rangle \langle A(\mathbf{s}_2) \rangle} \right], \quad (\text{A8})$$


giving us the relationship between the covariance function of  $\log A$  and the density power spectrum. In the following section we take the dust density model described in section 2 and test the assumption that  $\log A$  can be modelled as a Gaussian random field with the covariance function (A8) derived here.

## APPENDIX B: SIMULATIONS OF THE RELATIONSHIP BETWEEN THE CORRELATION FUNCTIONS OF $\log \rho$ AND $\log A$


### B1 Simulating Gaussian Random Fields

In order to test our method for mapping extinction we need to produce simulations of the ISM with realistically varying dust density. In order to do so we approximate  $\log(\rho)$  as a GRF which we then simulate, before exponentiating to obtain a simulation of  $\rho$ . This was the approach





**Figure B1.** Slices through 3 lognormal random field cubes produced by exponentiating GRFs. On the left a field simulated assuming  $\gamma = 3$ , in the centre  $\gamma = 3.5$  and on the right  $\gamma = 4$ .



**Figure B2.** A histogram of  $\log A$  obtained at a single position for 100,000 simulations. Overplotted in red is a Gaussian distribution with a mean of zero and unit variance.

adopted by Coles & Jones (1991), Elmegreen (2002) and Fischera & Dopita (2004). The later two seek to simulate an absorbing screen in the ISM using particular form of a Gaussian random field: fractional Brownian motion (fBm).

There are a number of methods for simulating Gaussian random fields, Saupe (1988) summarises several methods for producing fBm fields. A popular method, owing to its speed and ease of implementation, is the spectral synthesis method. This method relies on the fact that the power spectrum is the square of the Fourier transform of the field. As it uses Fourier transforms, this method must be used with care as this approach implicitly assumes that the field it is simulating is periodic in real space. However, as noted in section 2.2.3, the power spectra we employ are truncated at small  $\mathbf{k}$ , which helps disguise the periodic nature of the simulated fields. Additionally, the use of Fourier transforms means that it is only possible to directly simulate stationary fields.

In the spectral synthesis method one proceeds by first taking a cube of data points in  $\mathbf{k}$  space. Each point is then assigned a complex magnitude following the square root of the power spectrum assumed and a complex phase drawn randomly in the range 0 to  $2\pi$ . By then taking the Fourier transform of the cube, a data cube in real space is obtained where the data points follow a Gaussian random field with the desired power spectrum.

We would like to simulate a density field such that  $|\tilde{\rho}(\mathbf{k})|^2 \propto \mathbf{k}^{-\gamma}$ . Fischera & Dopita (2004) found that, for the range  $3 \leq \gamma < 4$ , a satisfactory density field can be produced by exponentiating a Gaussian random field for  $\log \rho$  which exhibits a power-law power spectrum  $|\log \tilde{\rho}(\mathbf{k})|^2 \propto \mathbf{k}^{-\beta}$ .


We can then modify the stationary dust density field by multiplying it by a slowly varying function that represents the variations in mean density, as discussed in section 2.2.3.

Following this groundwork, it is now possible to simulate dust density fields. Three example fields are shown in Fig. B1. Subsequently one can insert stars into the medium and integrate between the observer and the star to obtain the extinction to the star. We employ this scheme to produce the simulated data we study in section 4.

## B2 Comparing simulations to our assumed model

As discussed in section 2, we assume that  $\log A$  can be modelled as a Gaussian random field. If this is true then the pdf of the logarithm of extinction to a set position in space should be Gaussian. Here we test this assumption for the specific form of density fluctuations adopted in Appendix B1 above, in which we assumed that  $\log \rho$  was also a GRF. To do this we simulate 100,000 dust density field using the method described above. Specifically, we assume  $\gamma = 11/3$  and simulate fields 2048 pixels deep and 128 wide, with a scale length of 16 pixels. We then integrate along a distance of 480 pixels, which is equivalent to 3 kpc if the outer scale length is 100 pc, to obtain a simulated extinction. Fig. B2, shows the histogram of the simulated  $\log A$ , which is indeed consistent with being Gaussian.





**Figure B3.** The calculated covariance of extinction from 10,000 simulated fields. Plotted here is the dependence of covariance on the length of the shorter sightline,  $\min(s_1, s_2)$ , and the difference in length between the two sightlines,  $\Delta s = |s_1 - s_2|$ .

The condition for a field to be a GRF is that it satisfies the condition (7): from this it should be clear that, in general, if  $\log p$  satisfies this condition, it does not follow that  $\log A$  does and vice versa. However, it has been realized in a variety of contexts, such as wireless communications (e.g., Mehta et al. 2007) and finance (e.g., Ju 2002), that a sum of lognormal random variables can be well approximated by a single lognormal random variable. Barakat (1976) describes how, due to their non-zero skew, a sum of lognormal random variables converges only slowly onto a Gaussian distribution under the central limit theorem and how, prior to convergence, the sum is better approximated by a lognormal distribution.

In the previous appendix we derived a form for the covariance function of  $A$ ,  $\text{Cov}[A(\mathbf{r}_1), A(\mathbf{r}_2)]$ , to two positions. We showed that  $\text{Cov}[A(\mathbf{r}_1), A(\mathbf{r}_2)]$  depends on the angular separation of the positions and the distance to the nearer position only. Consequently we have a non-stationary covariance function, where we assume that  $\text{Cov}[A(\mathbf{r}_1), A(\mathbf{r}_2)]$  is independent of the difference between the two distances. We employ a subset of 10,000 of the 100,000 simulations produced above to study the covariance function. We consider a sub region of 1760 by 64 pixels, dropping the rest of the field to avoid the affects of periodicity in the simulations. We then calculate the covariance for every remaining pixel pair. In Fig. B3 we show that  $\text{Cov}[A(\mathbf{r}_1), A(\mathbf{r}_2)]$  is indeed well approximated as a linear function of the shorter distance. Moreover, we show that the covariance is sensitive to the difference in distances only in the case that the two sightlines have approximately equal length. The reason this arises is that any sections of the longer sightline that are substantially more than one scale length beyond the end of the shorter sightline will be essentially independent of the shorter sightline, whilst those sections of the longer sightline within a few scale lengths of the end of the shorter one will comprise only a small proportion of the longer sightline.